The AMD-760™ MP Platform

And Then There Were Two…

Jessie J. Johnson
ADVANCED MICRO DEVICES, INC.
One AMD Place
Sunnyvale, CA 94088

Version 0.2

Table of Contents:

Introduction:

High-performance…scalability…manageability, these are key terms that describe a class of power platforms intended for mission-critical applications required of servers and workstations. The AMD-760™ MP platform is a high-performance, two-way multiprocessor (MP) system solution designed for server- and workstation-class applications. Consisting of the award-winning AMD Athlon™ processor and the AMD-760 MP chipset, this solution offers scalable processing capability, high-bandwidth memory and I/O performance, and sophisticated system management capability to support a wide range of computing infrastructures. This white paper describes the AMD-760 MP platform, architecture, and underlying technologies.

AMD-760™ MP Platform Overview:

The AMD-760 MP platform addresses the server and workstation sectors by offering the following high-level features:

- Uniprocessor and two-way symmetric multiprocessing capability, supporting AMD Athlon class processors
- Dual point-to-point 266MHz AMD Athlon system buses designed to support up to 2.1GB/s transfer rate per system bus
- 266MHz DDR (Double Data Rate) memory interface supporting up to 4GB of memory space using registered PC2100 DIMMs; ECC (Error Correcting Code) memory is also supported
- AGP-4x graphics interface, backwards-compatible with AGP-1x and 2x modes
- 33MHz/32-bit/64-bit PCI 2.2-compliant PCI bus interface
- EIDE storage controller subsystem supporting ATA-33/66/100 MB/s data rates
- System management functions

As shown in Figure 1, these features are packaged in a two-chip core logic solution consisting of the AMD-762™ system controller (Northbridge) and the AMD-766™ peripheral bus controller (Southbridge). When implemented with AMD Athlon processor technology, these elements combine to deliver outstanding performance to server-and workstation-class systems.

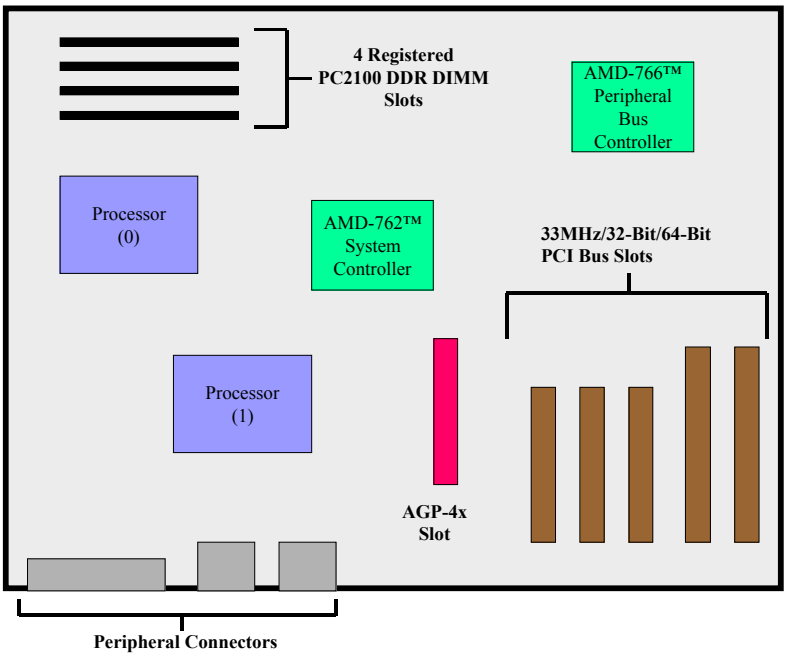

Figure 1: **AMD-760™ MP Motherboard Platform**

Platform Architecture:

Servers and workstations leverage similar underlying technologies as fundamental building blocks. Both classes of systems require high-performance computing capability, high-capacity memory elements, and robust I/O (Input/Output) subsystems. These platforms, however, begin to diverge in the areas of scalability, graphics, and I/O control. Workstations, for example, require high-end graphics, mid-level storage capabilities, and limited I/O expansion. On the other hand, servers do not require high-end graphics, but instead require "scalable" high-performance computing capabilities, massive high-performance storage engines, and an I/O subsystem that is robust enough to concurrently support a multitude of high-speed networking and storage subsystems.

Early in the design process, AMD engineers faced the extreme challenge of developing an architecture that is powerful, robust, and cost-effective, yet flexible enough to support the diverging demands of both server- and workstation-class systems. Their answer: the AMD-760 MP platform.

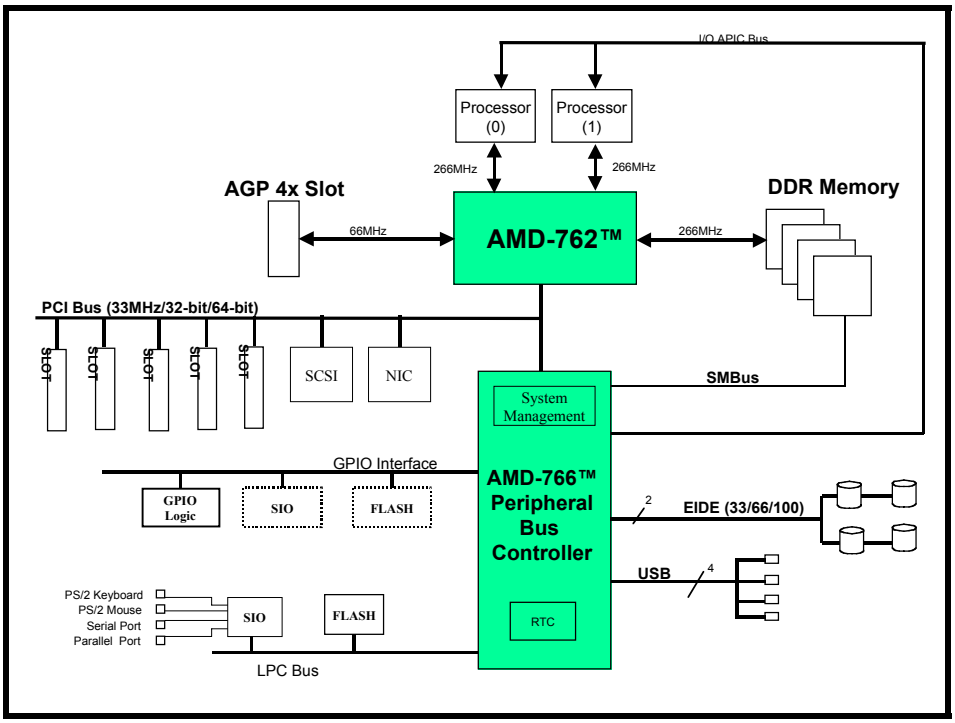

Figure 2: **AMD-760™ MP Platform High-Level Architecture**

As shown in Figure 2, the AMD-760 MP platform employs PGA-socketed AMD Athlon processor technology as the high-performance system engine(s) of choice. The AMD-760 MP design supports a single processor for workstation-class systems or low-end servers, leaving a second available processor socket for future expansion. Customers will enjoy the security of knowing their original investment is protected with a system that can be scaled as their computing demands grow. Customers can upwardly scale the processing capability of their system in one of several ways:

1. Replacing the original single processor with a faster or improved AMD Athlon processor

2. Adding a second AMD Athlon processor (of the same clock speed as the first processor), thus enabling symmetric multiprocessing and significantly boosting system performance
3. Replacing the original single processor with a faster AMD Athlon processor and adding a second processor of equivalent speed

In high-end workstation and mid-range server applications, where maximum performance is required from the onset, systems can be shipped with two processors, immediately enabling the benefits of symmetric multiprocessing.

The system logic is partitioned as follows:

The AMD-762™ system controller houses the high-speed elements critical to system performance. Contained in the AMD-762 controller are the following subsystems:

- Dual point-to-point 266MHz AMD Athlon system bus interfaces, supporting up two processors
- 266MHz DDR memory interface, supporting up to 4GB of memory space using PC2100 registered DDR memory DIMMs; ECC memory is also supported
- AGP-4x graphics interface
- 33MHz/32-bit/64-bit PCI bus interface
- 949-pin plastic BGA (Ball-Grid Array) package
- 2.5V core

For detailed descriptions, refer to the **AMD-762™ System Controller Data Sheet**.

The **AMD-766™ peripheral bus controller** compliments the AMD-762 system controller by offering a robust I/O subsystem along with sophisticated system management capability. Embedded are the following features:

- 33MHz/32-bit PCI bus interface
- ATA-33/66/100 EIDE interface
- Four-port OHCI USB host controller
- LPC (Low Pin Count) bus
- System management interface
- SMBus interface
- Real time clock (RTC)
- 256 bytes CMOS memory
- I/O APIC interrupt controller
- GPIO interface
- 272-pin BGA package
- 3.3V core and output drivers; 5V tolerant input buffers

For detailed descriptions, refer to the **AMD-766™ Peripheral Bus Controller Datasheet.**

Dual AMD Athlon™ System Bus:

Server- and workstation-class platforms require high-performance and scalable compute capability. A key element impacting these requirements is the system bus linking the processor(s) to the system logic. Throughout the industry, this bus is generally termed the "front-side bus."

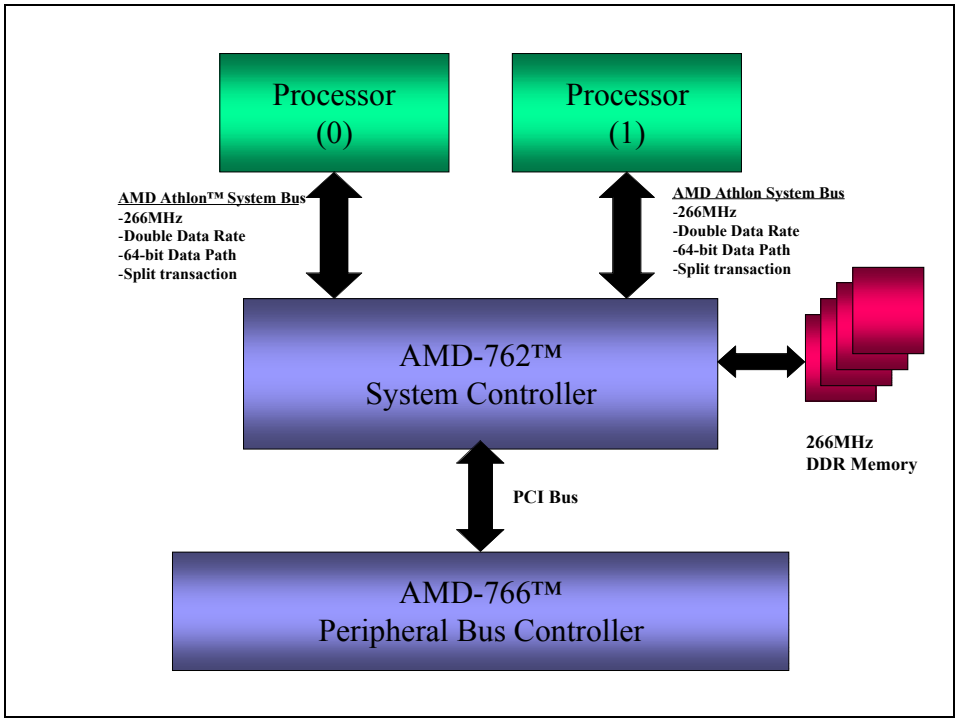

Figure 3: **Dual Point-to-Point AMD Athlon™ System Bus Architecture**

As shown in Figure 3, the AMD-760 MP platform employs the AMD Athlon system bus as the high-speed interface between the processor(s) and system elements. Optimized for server and workstation applications, the AMD Athlon system bus offers the following attributes:

- 266MHz[1] operation, designed to deliver a peak throughput of 2.1GB/s per system bus
- Split-transaction architecture
- Cache-coherency protocol
- ECC capability
- 64-bit data path
- Packetized request transactions

[1] The 266MHz rate refers to a physical clock operating at 133MHz in which information is transferred on each clock edge. Hence, this is calculated as (133MHz clock) x (2 transfers/clock) = 266M transfers/sec = 266MHz.

To achieve symmetric multiprocessing capability, AMD chose a dual point-to-point independent bus topology as opposed to the shared bus architecture depicted in Figure 4. The dual point-to-point architecture offers the following advantages over the shared bus architecture:

- Processors on the AMD Athlon system bus can concurrently burst information to and from the system. The shared bus architecture allows only one processor at a time to transfer data.
- Latency is reduced as processors on the AMD Athlon system bus will not have to arbitrate for system bus control.
- Signal loading, integrity, and termination are much simpler from a design perspective with only one processor per system bus.

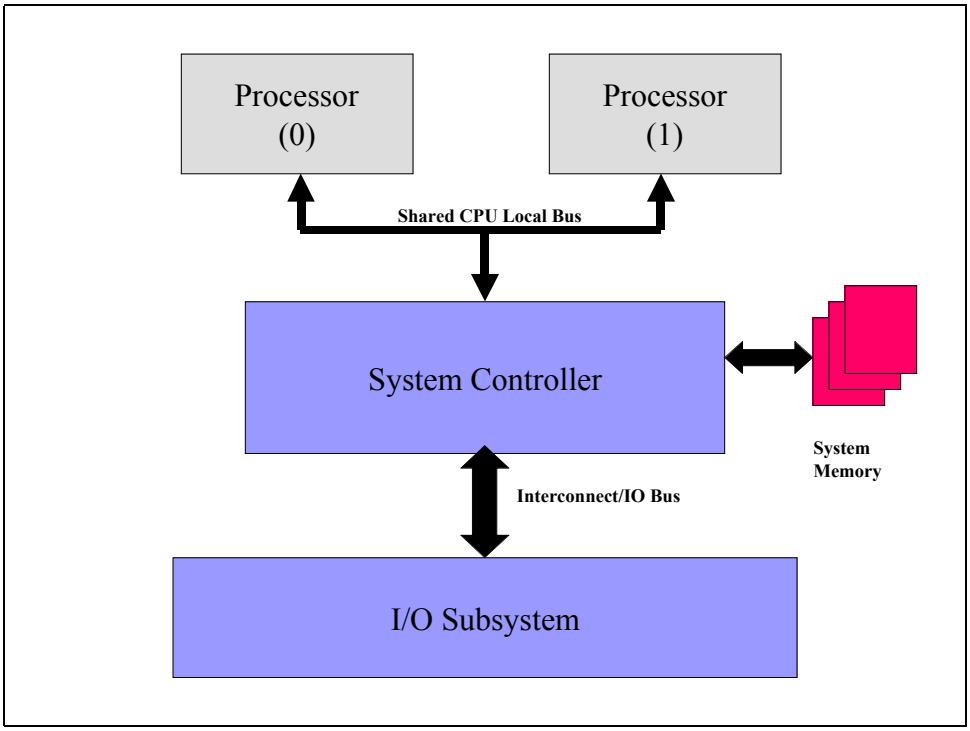

Figure 4: **Generic Shared CPU Local Bus Architecture**

Split-transaction bus architectures have existed in mini, mid-range, and mainframe computer systems for many years. Only recently have these architectures moved into the desktop market. In essence, split-transaction architecture is a bus transaction technique in which serialized transactions are decoupled and allowed to execute concurrently, or in an overlapping fashion. This technique leverages the latencies in the system to increase overall system performance.

This architecture is different from traditional microprocessor bus architectures in which the address, control, and data busses operate collectively as a single pipe (see Figure 5). The traditional architecture allows only a single transaction to occur at a time, thus creating a serializing affect.

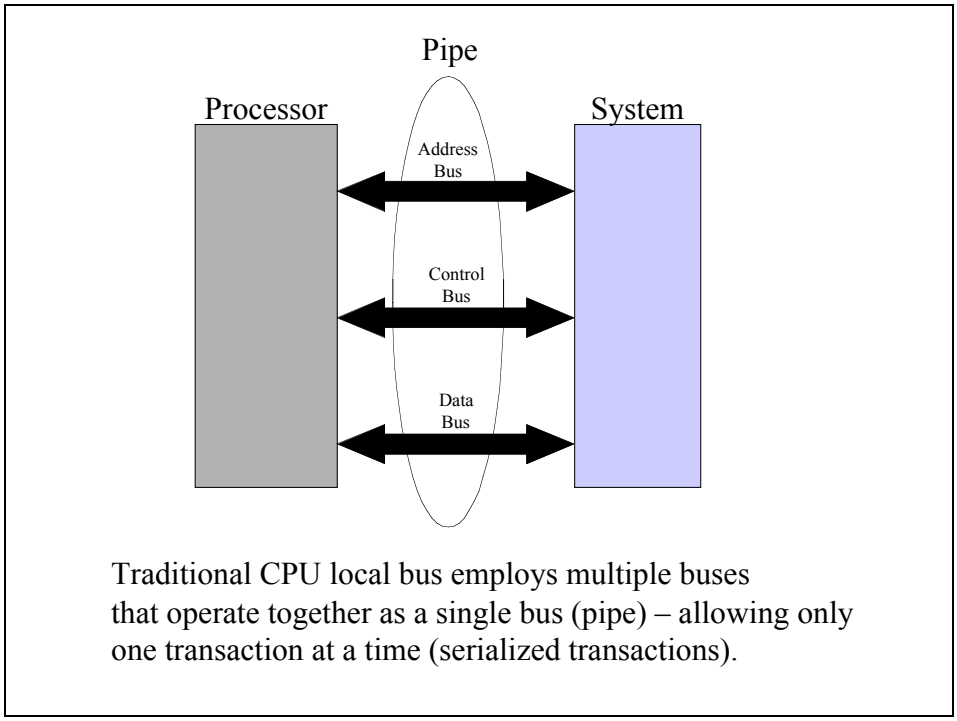

Figure 5: **Traditional Microprocessor CPU Local Bus**

Figure 6 is an example that illustrates the serialization nature of the non-split transaction type buses. Bus transactions can be decomposed into an address/control phase and a data phase. In the case where the processor is reading data from the system, during the address/control phase, address and control information is broadcast to the system. The processor must wait for the system to respond with valid data before a new transaction can be launched.[2] The length of the delay depends on the type of device (memory, disk drive, etc.) being accessed, its associated latency, and its readiness to deliver data. In Figure 6, note how the data phase times can vary from transaction to transaction. In this example, it takes 15 units of time to complete three transactions.

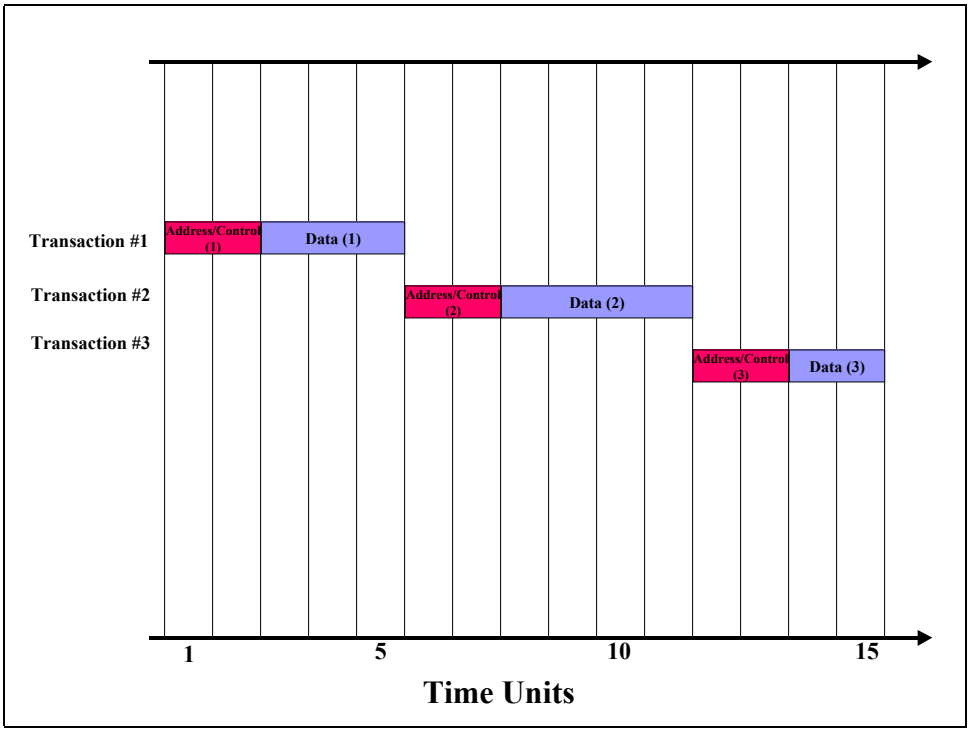

Figure 6: **Example of Serialized Transaction Nature of Non-Split Transaction Bus Architectures**

Split-transaction architecture improves performance by taking advantage of the data phase latencies, and allowing other transactions to launch as the processor (or system) is waiting for the current data phase to complete. Hence, transactions overlap, concurrency is achieved, and system performance is increased.

[2] Techniques such as pipelining and delayed transaction are used to circumvent this type of issue; however, these techniques will be ignored in this white paper to simplify the discussion.

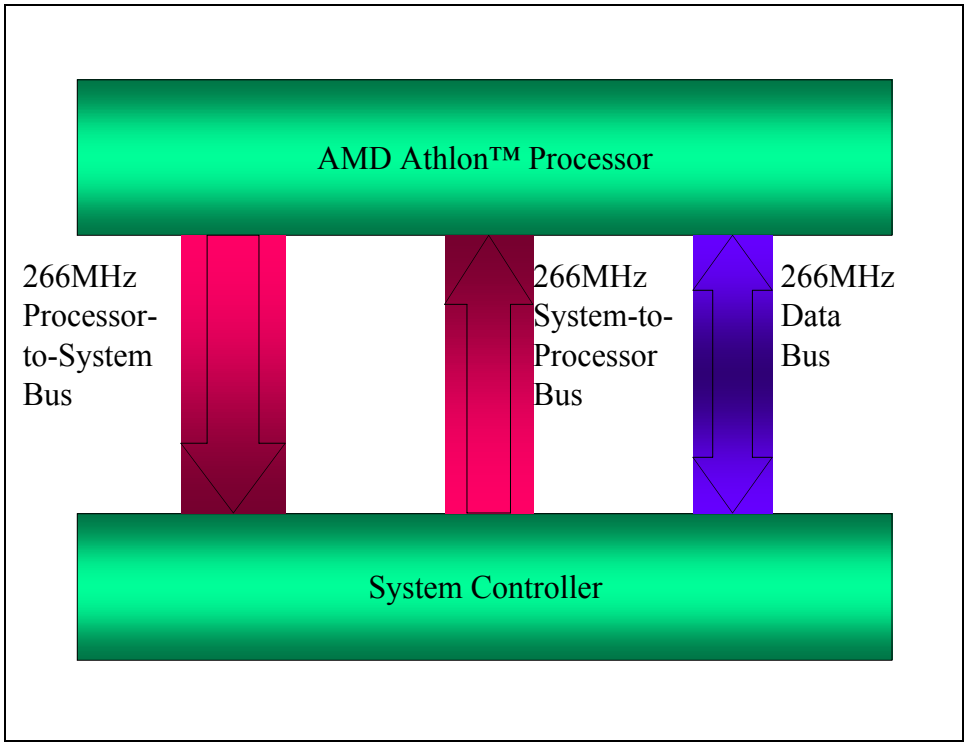

Figure 7: **AMD Athlon™ System Bus Components**

As shown in Figure 7, the AMD Athlon system bus is composed of three separate buses that operate independently and concurrently:

1. Processor-to-system bus
2. System-to-processor bus
3. Data bus

Note that each bus runs at 266MHz[3].

The **processor-to-system bus** is the channel in which the processor issues requests/commands (memory read, memory write, etc.) to the system. This is a uni-directional bus controlled only by the processor. Data is not transferred on this bus; only packets containing command, address, and other information are transferred. Data information is transferred over the data bus.

[3] The 266MHz rate refers to a physical clock operating at 133MHz in which information is transferred on each clock edge. Hence, this is calculated as (133MHz clock) x (2 transfers/clock) = 266M transfers/sec = 266MHz.

The **system-to-processor bus** is the channel in which the system issues requests/commands to the processor. This is also a uni-directional bus controlled only by the system controller. Similar to the processor-to-system bus, only packets—request and command information—are transferred on this bus. Data information is transferred over the data bus.

The **data bus** is a bi-directional bus used to transfer data packets between the processor and system elements in response to requests from their respective buses. Each data packet contains an ID tag for association with the corresponding processor-to-system or system-to-processor request.

NOTE: Refer to the **AMD Athlon™ System Bus Specification** for more details on bus architecture and operation.

As shown in Figure 8, the three independent, high-performance channels (and support logic) enable split-transaction capability. Processor requests (PR) and system requests (SR) can be sent concurrently and independently. System data (SD) and processor data (PD) can be returned to the requestor immediately or at a later time, depending on availability. Ordering is mitigated by instantiating a unique ID associating the request with the corresponding data.

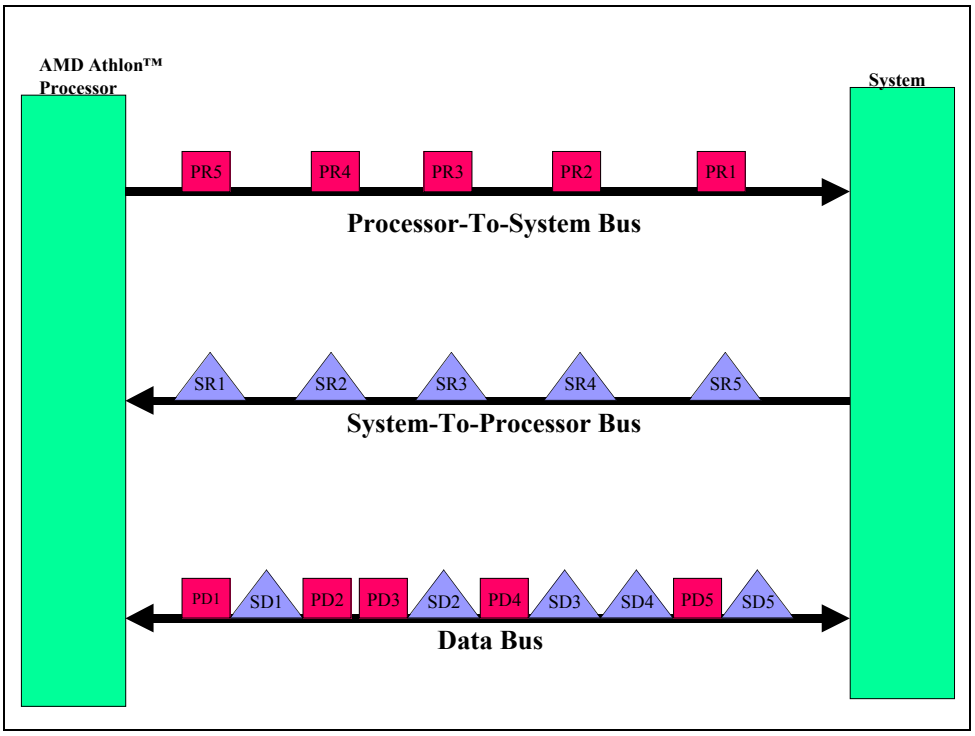

Figure 8: **AMD Athlon™ System Bus Split Transaction Example**

The net affect of split-transaction capability is improved performance. As shown in Figure 9, concurrency is achieved by decoupling the processor-to-system transactions from the system-to-processor transactions, and by decoupling the data response from the associated request. Compared to the non-split transaction example shown in Figure 6, the time to complete three transactions has been reduced from 15 time units to 11 time units[4].

[4] These performance numbers are arbitrary examples used for concept illustration only.

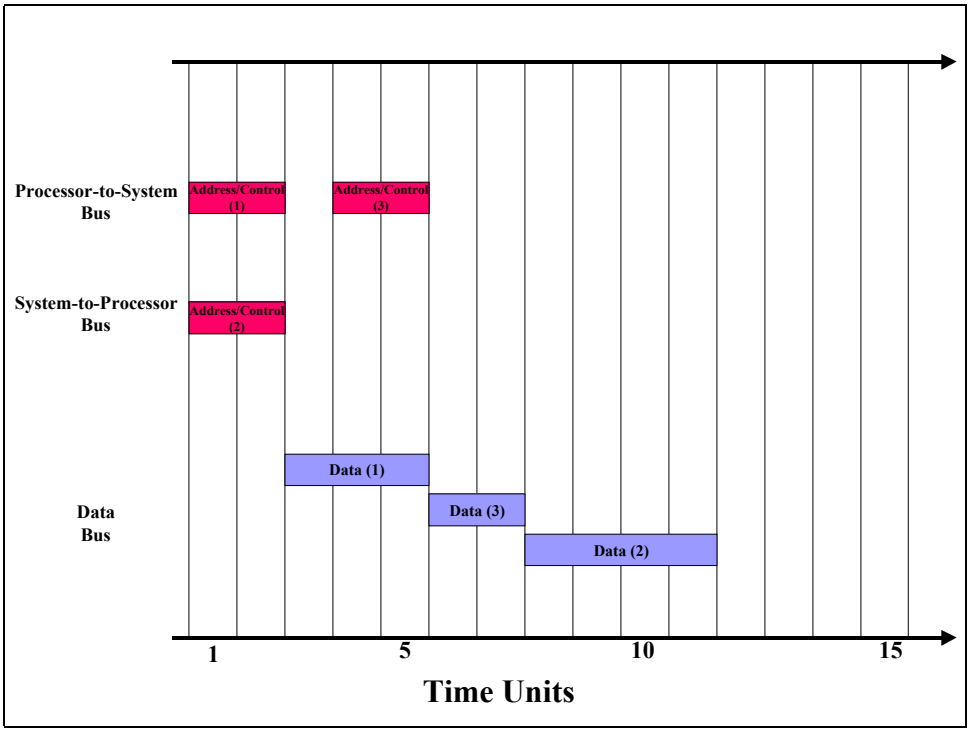

Figure 9: **Example Illustrating the Concurrency Offered by Split-Transaction Bus Architecture**

The AMD Athlon system bus also supports error checking and correcting across the system bus to evaluate data integrity.

NOTE: When running two processors, both processors must run at the same clock frequencies. Also, the AMD Athlon system bus and DDR memory subsystem are locked in frequency. Whether running one processor or two processors, the AMD Athlon system bus and DDR memory subsystem must operate at the same speed.

In summary, the AMD Athlon system bus offers high-speed communication between the processor and the system. The dual point-to-point multiprocessor topology, split-transaction capability, and error checking and correcting features all combine to deliver outstanding performance and protection for mission-critical applications.

DDR Memory Subsystem:

The AMD-760 MP platform implements the latest evolution in high-performance memory technology—Double Data Rate (DDR) SDRAM. DDR memory is natural extension of current PC100/PC133 SDRAM technology. However, DDR offers higher performance at a competitive cost.

Similar to the AMD Athlon system bus, the DDR memory subsystem is designed to deliver a peak transfer rate of up to 2.1GB/s. To satisfy the large memory requirements of server- and workstation-class systems, the AMD-762 system controller supports up to 4GB of memory space and provides interfacing for four registered DIMM slots.

NOTE: Unbuffered memory DIMMs are not supported.

In applications requiring minimal fault tolerant capability, the AMD-760 MP DDR memory subsystem supports ECC memory and provides detection and correction of single-bit errors.

AGP-4x Graphics Subsystem:

The graphics subsystem is a critical element of high-performance workstations. The AMD-760 MP platform implements AGP (Accelerated Graphics Port) technology as its graphics subsystem of choice. The intent of the AMD-760 MP platform is to increase overall system performance by eliminating data movement bottlenecks, allowing a more efficient match to the compute performance offered by faster AMD processors. Increasing performance without impacting system cost is a difficult design challenge; however, AGP graphics technology offers the perfect balance to satisfy even the most extreme user.

The AMD-762 system controller is equipped with an AGP-4x graphics interface designed to provide powerful graphics capability to the workstation desktop. Optimized to run concurrently with the AMD Athlon system bus and DDR memory interface, high-end graphics adapters can truly utilize up to 1GB/s of bandwidth potential offered by the AGP-4x subsystem.

Server applications typically do not require high-end graphics capability. The AMD-762 system controller's AGP-4x interface is backwards-compatible with lower performance AGP-1x and AGP-2x modes. This feature allows system designers the flexibility to leverage lower-cost AGP graphics adapters in their system solutions, thus reducing overall platform cost.

ATA-100 Storage Subsystem:

Due to the physical/mechanical nature of hard drives, the storage subsystem is one of the slowest elements within all computer systems. This has a significant impact on overall system performance, since all operating system and software applications are initially loaded from the hard drive. Essentially, the storage subsystem "data pipe" becomes a critical bottleneck in the system.

To alleviate this performance issue, the AMD-760 MP platform implements a storage subsystem controller that offers increased data transfer rates between the EIDE hard drive controller and the actual storage device (hard drive, CD/CD-R/CDRW, DVD, etc.). As shown in Figure 10, the AMD-766 peripheral bus controller embeds an EIDE storage controller that offers data transfer rates of 33MB/s, 66MB/s, and 100MB/s and are compliant with the ATA/UDMA-33/66/100 standards (shown in Table 1).

Although 5400 RPM drives are supported, optimal EIDE storage subsystem performance is obtained when the ATA-100 mode is used in conjunction with hard drives that operate at spindle speeds of 7200 RPM or higher.

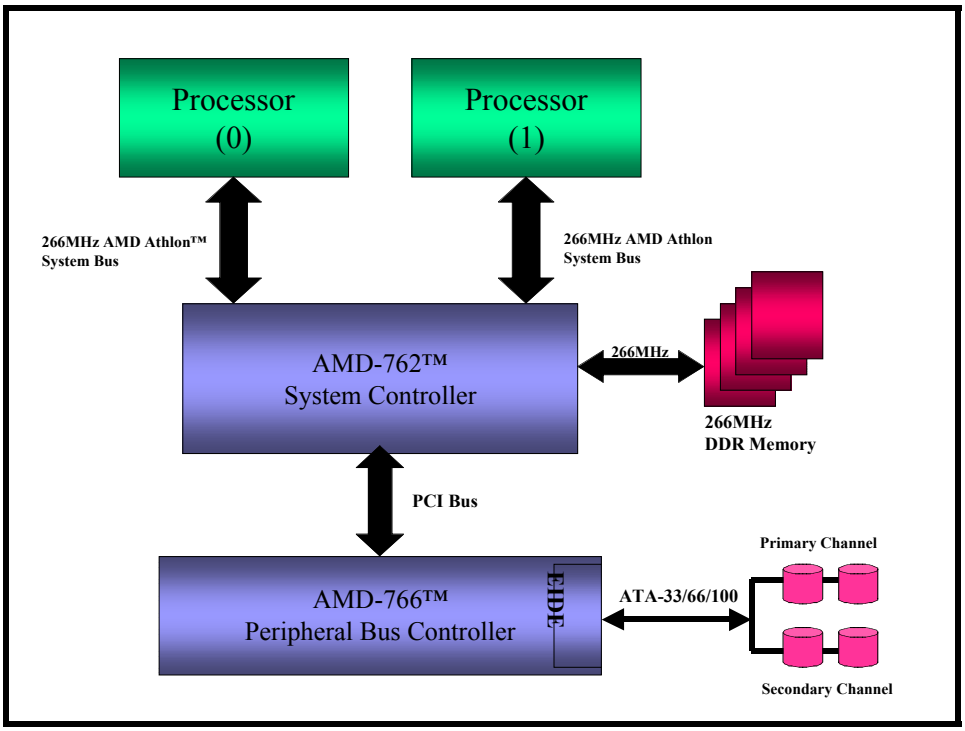

Figure 10: **The AMD-760™ MP Chipset EIDE Controller Supporting ATA-33/66/100 MB/s Speeds**

Table 1: **AMD-760™ MP Chipset EIDE Interface Speeds**

ATA Mode	Data Transfer Rate
ATA/UDMA 33	33MB/sec.
ATA/UDMA 66	66MB/sec.
ATA/UDMA 100	100MB/sec.

PCI Bus:

The AMD-760 MP chipset utilizes the PCI 2.2-compliant 33MHz PCI bus as the I/O backbone of the system. As shown in Figure 11, this bus serves several purposes:

1. Interconnect between the core logic elements
2. I/O expansion via PCI slots
3. Custom functionality via embedded motherboard devices

The AMD-760 MP chipset has hardware arbitration capability to support up to seven PCI devices[5]. These seven devices can be a mixture of slots and on-board motherboard components[6].

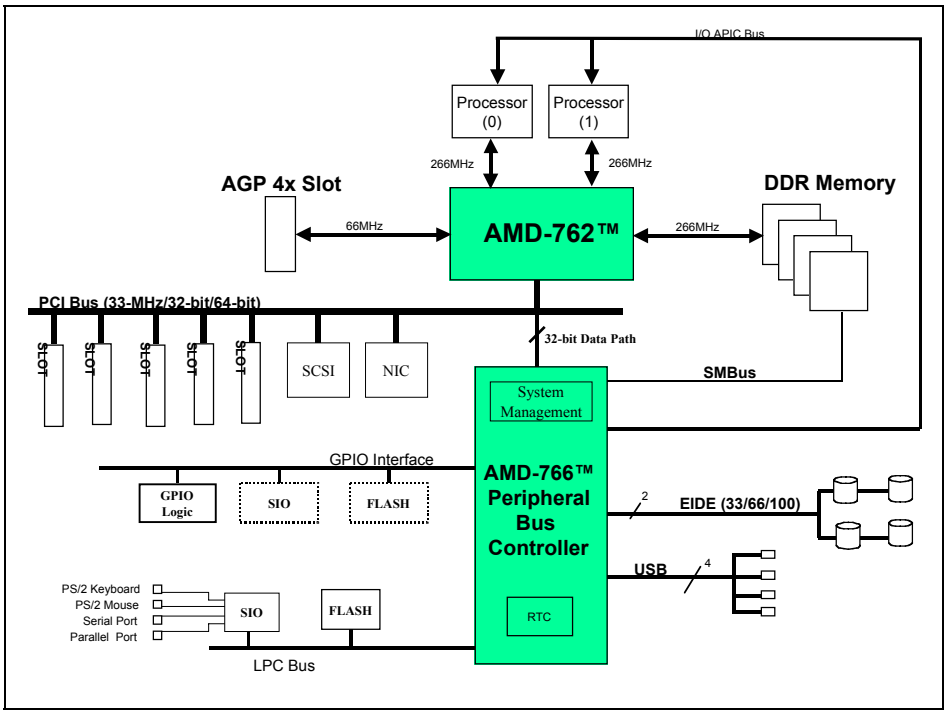

Figure 11: **AMD-760™ MP Architecture Illustrating PCI Bus Connectivity**

To allow flexibility in cost/performance, the AMD-760 MP chipset supports two performance modes at 33MHz—32-bit data path and 64-bit data path. As shown in Table 2, PCI devices implementing a 32-bit data path can approach peak transfer rates of 132MB/sec., while higher-performance devices/adapters implementing 64-bit data paths can approach peak transfer rates of 264MB/sec. Both can be supported in the same system design.

[5] REQ/GNT#[6:0] pairs are allocated for PCI slots and on-board devices. A separate pair of signals (SBREQ/SBGNT#) is dedicated to the AMD-766 peripheral bus controller.

[6] AMD designs and tests its reference design motherboards to support 5 PCI slots. If more slots or additional motherboard devices are desired, the system designer must perform the appropriate analysis and testing to ensure signal integrity and proper operation.

Table 2: **PCI Bus 33MHz Peak Data Transfer Rates with 32-Bit and 64-Bit Data Paths**

Data Path Width	Peak Bandwidth
33MHz, 32-bit data path	132MB/sec.
33MHz, 64-bit data path	264MB/sec.

These performance modes allow the designer flexibility to balance cost and performance. For example, in workstation applications, inexpensive 32-bit network adapters can be used for connectivity, while higher-performance 64-bit adapters can be used to enable a fast SCSI storage subsystem. In server applications, high-performance 64-bit adapters can be used to enable a massive RAID storage subsystem or to accommodate several high-speed multi-port network adapters.

NOTE: The AMD-766 peripheral bus controller implements a 32-bit PCI bus interface.

Appendix:

Glossary of Terms:

AGP—Accelerated Graphics Port
BGA—Ball Grid Array
DDR—Double Data Rate
ECC—Error Correcting Code
I/O—Input/Output
MP—Multiprocessor
SBA—Side-Band Addressing

DDR PC2100/PC1600 Definitions:

PC2100 modules are DDR memory modules designed to support a peak data rate of 2100MB/sec. This data rate is calculated as follows:

- PC2100 peak data rate = (133MHz clock) * (2 transfers/clock) * (8 bytes/transfer)
- PC2100 peak data rate = 2.128GB/sec.
- PC2100 peak data rate = 2.1GB/sec. (rounded)
- PC2100 peak data rate = 2100 MB/sec.

PC1600 modules are DDR memory modules designed to support a peak data rate of 1600MB/sec. This data rate is calculated as follows:

- PC1600 peak data rate = (100MHz clock) * (2 transfers/clock) * (8 bytes/transfer)
- PC1600 peak data rate = 1.6GB/sec.
- PC1600 peak data rate = 1600MB/sec.

AGP Graphics Background:

To understand the advantages and benefits of AGP graphics, it is necessary to understand the issues that AGP technology has resolved. Figure 12 shows an architectural diagram of a generic PCI bus-based graphics subsystem. In this architecture, the graphics subsystem resides on the PCI bus. Note that the PCI bus graphics adapter embeds its own local memory on the adapter card. Although this architecture performed well in its time frame, several issues arose that motivated the need for AGP:

1. Upgrading graphics memory was expensive, as memory modules must be added to the graphics card, or the graphics card must be replaced entirely.
2. Since some graphics data (such as textures and other information) are stored in main memory, the PCI-based graphics card must access main memory via the PCI bus. These accesses may occur frequently, particularly if the graphics adapter has a small amount of local memory. Unfortunately, the graphics card must compete with other PCI bus peripherals for PCI bus bandwidth.
3. If the graphics adapter must make frequent PCI bus accesses, other PCI bus peripherals may become starved for PCI bus bandwidth.

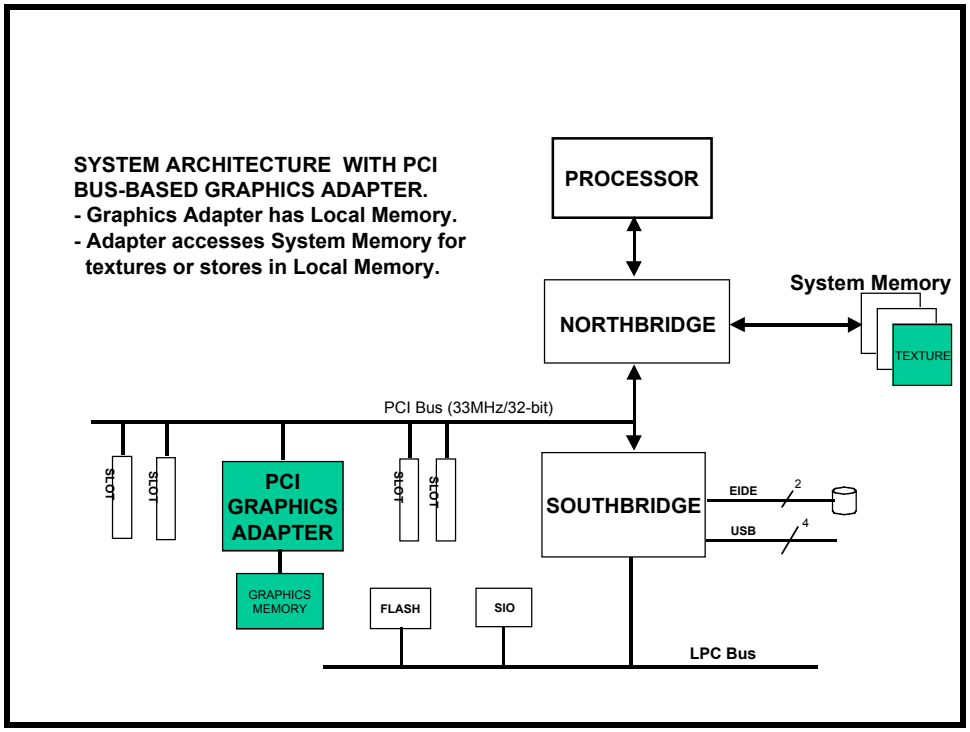

Figure 12: **Older System Architecture Showing PCI Bus-Based Graphics**

Figure 13 illustrates how AGP technology elegantly resolves the issues facing PCI bus graphics architecture. The AMD-762 system controller (the Northbridge element of the AMD-760 MP chipset) embeds the system's AGP graphics interface. The AGP interface utilizes the 66MHz PCI bus protocol in tandem with a side-band addressing (SBA) bus for concurrent posting of commands from the graphics card to the AGP logic embedded in the Northbridge (See Figure 14). The Northbridge embeds read/write and command queues (buffers) to allow full-speed data and command transport between the AGP device and AMD-762 system controller, and concurrent full-speed data transport between the AMD-762 system controller and the DDR memory subsystem.

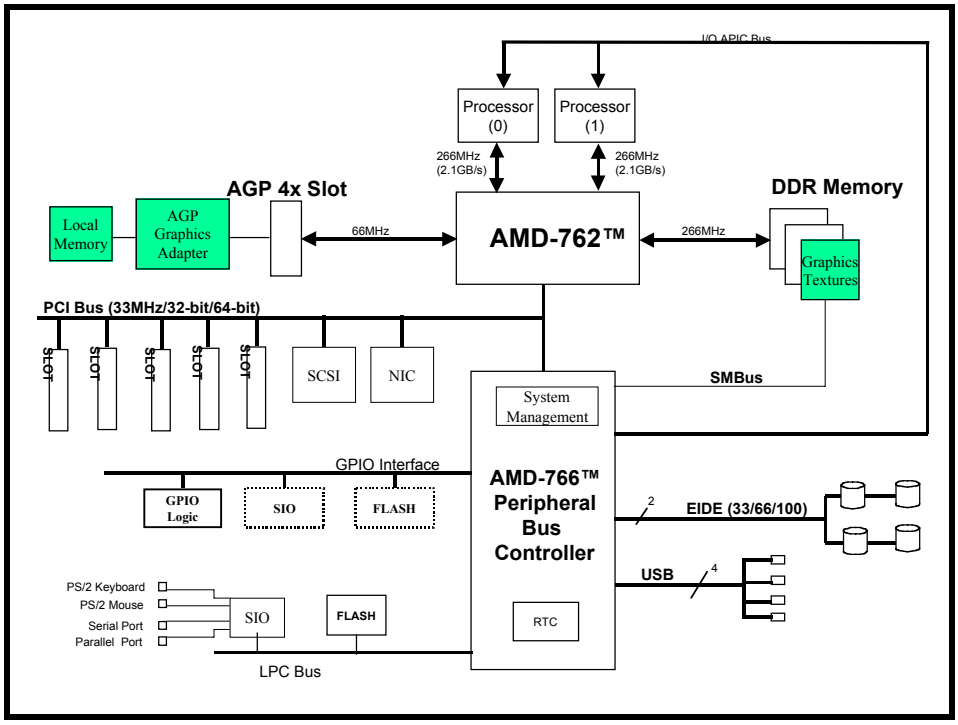

Figure 13: **AMD-762™ System Controller AGP Graphics Architecture Showing Use of System Memory for Graphics Operations**

The architectures shown in Figure 13 and Figure 14 produce the following benefits:

- The native architecture of AGP graphics subsystem (66MHz PCI bus interface with side-band addressing, and embedded Northbridge AGP logic) offers significant raw performance improvement over PCI bus based graphics subsystems.
- The AGP architecture allows the AGP graphics subsystem to view and use main memory just like it's own local memory—meaning that the AGP graphics card shares system memory. The AGP graphics card cannot distinguish between system memory and local memory, as it all appears as local memory. To the end-user, graphics performance can be enhanced by increasing system memory (inexpensive), rather than by adding expensive graphics memory.

- The graphics subsystem no longer has to compete for PCI bus bandwidth to access data in system memory. This benefit allows the graphics subsystem to run at full speed with minimal interruption from other components in the system. It also increases system concurrency—meaning that the processor, AGP graphics subsystem, and PCI bus device can run independently and concurrently, thus increasing system performance.
- PCI bus devices no longer have to compete with the graphics adapter for PCI bus bandwidth. PCI bus availability has been increased with the removal of the graphics subsystem from the PCI bus.

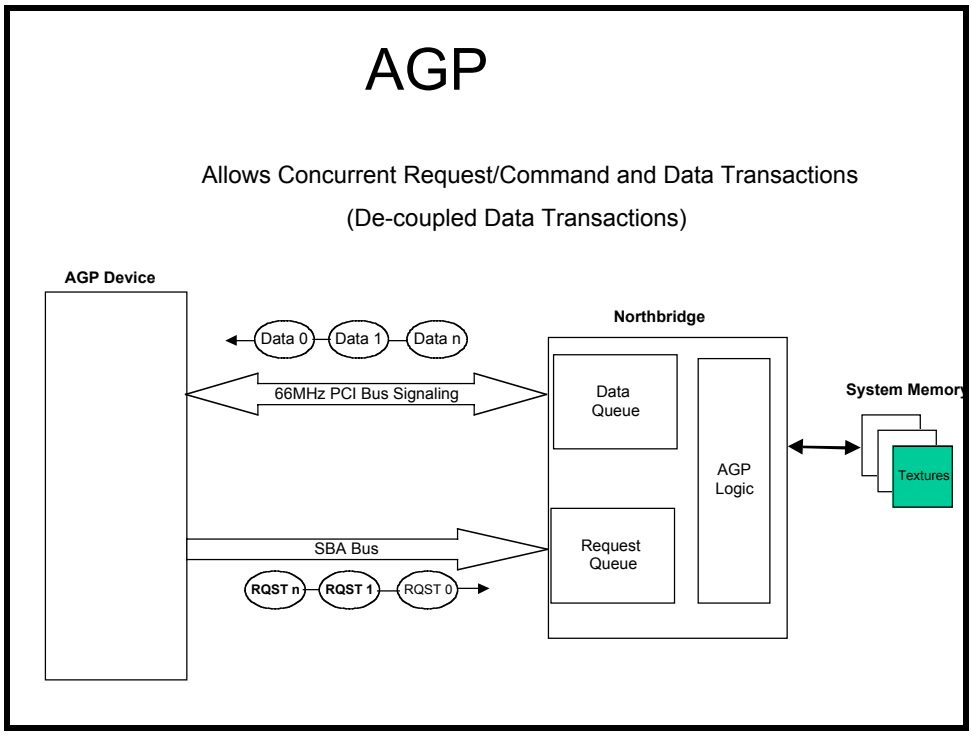

Figure 14: **High-Level AGP Interface Diagram Showing Bus Architecture and Embedded Northbridge Components**

Over time, the AGP graphics subsystem has scaled to increasing levels of performance. As shown in Table 4, there are several modes (data transfer rates) that have evolved over time. Analogous to gears on a sports car transmission, first gear is the original AGP-1x mode, offering a data transfer rate of up to 264MB/sec. Second gear is AGP-2x mode, doubling the data transfer rate to up to 528MB/sec. Finally, third gear is AGP-4x, offering the highest performance data transfer rate of up to 1GB/sec. (The notation 2x and 4x

are relative the original AGP-1x mode). In the future, higher performance modes (overdrive and turbo gears, perhaps) may be added.

Table 3: **AGP Modes and Corresponding Peak Bandwidths.**

AGP Graphics Mode	Peak Bandwidth (data transfer rate)
AGP-1x	Up to 264MB/sec.
AGP-2x	Up to 528MB/sec.
AGP-4x	Up to 1GB/sec.

NOTE: The AMD-760 MP chipset is designed to support all modes shown in Table 4.

AMD Overview:

AMD is a global supplier of integrated circuits for the personal and networked computer and communications markets with manufacturing facilities in the United States, Europe, and Asia. AMD produces microprocessors, flash memory devices, and support circuitry for communications and networking applications. Founded in 1969 and based in Sunnyvale, California, AMD had revenues of $4.6 billion in 2000. (NYSE: AMD).